REPORT DOCUMENTATION PAGE
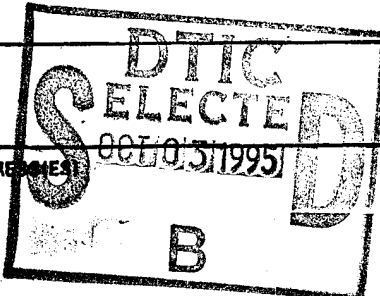
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of the collection of information, including suggestions for reducing the burden to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302 and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503

1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE September 5, 1995	3. REPORT TYPE AND DATES COVERED Final Technical Report 6/21/89 to 7/20/95

4. TITLE AND SUBTITLE OF REPORT

Reminding-based Learning

6. AUTHOR(S)

Brian H. Ross

5. FUNDING NUMBERS

Prog. element: 61102F
Proj.: 2313
TAsk: ~~A4~~ B3

DTIC SELECTED Oct 03 1995 B

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

University of Illinois
506 S. Wright Street
Urbana, IL 61801

8. PERFORMING ORGANIZATION REPORT NUMBER:

AFOSR-TR-95
0571

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

Air Force Office of Scientific Research/NL
Bolling AFB, DC 20332-6448

Attn: Dr. John Tangney (202)767-8075

NCY

11. SUPPLEMENTARY NOTES:

DISTRIBUTION STATEMENT A

Approved for public release;
Distribution Unlimited

12a. DISTRIBUTION AVAILABILITY STATEMENT Unlimited	12b. DISTRIBUTION CODE

19951002 008

13. ABSTRACT (Maximum 200 words)

When learning new cognitive skills involving problem solving, novices are often reminded of earlier problems. This project examined this common means of learning from remindings. First, the representation of the resulting generalization was investigated. Generalizations from earlier problems may be both selective (only some parts are included in the generalization) and conservative (some superficial aspects are included). The studies found evidence for these characteristics and showed how such generalization may be tied to the use. Second, these remindings may provide a means of becoming more expert in a problem solving domain. Experiments show that even highly experienced solvers rely upon superficial similarities that are predictive of the problem type. Third, an examination of remindings in everyday learning situations extended the findings and better tested some theoretical ideas. The overall results of this project provide a clearer understanding of reminding-based learning and relate it to work on expertise, categorization, and schema acquisition.

14. SUBJECT TERMS Learning, remindings, categorization, problem solving	15. NUMBER OF PAGES: 14
	16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT: unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT unclassified	20. LIMITATION OF ABSTRACT

FINAL TECHNICAL REPORT: REMINDING-BASED LEARNING

PI: Brian H. Ross, University of Illinois

AFOSR 89-0447

SUMMARY

When learning new cognitive skills involving problem solving, novices are often reminded of earlier problems. This project examined this common means of learning from remindings. First, the representation of the resulting generalization was investigated. Generalizations from earlier problems may be both selective (only some parts are included in the generalization) and conservative (some superficial aspects are included). The studies found evidence for these characteristics and showed how such generalization may be tied to the use. Second, these remindings may provide a means of becoming more expert in a problem solving domain. Experiments show that even highly experienced solvers rely upon superficial similarities that are predictive of the problem type. Third, an examination of remindings in everyday learning situations extended the findings and better tested some theoretical ideas. The overall results of this project provide a clearer understanding of reminding-based learning and relate it to work on expertise, categorization, and schema acquisition.

PUBLICATIONS ACKNOWLEDGING AFOSR SUPPORT

Ross, B. H., Perkins, S. & Tenpenny, P. L. (1990) Reminding-based category learning. Cognitive Psychology, 22.

Ross, B.H. (1990). The access and use of relevant knowledge: A specific case and general issues. In R. Freedle (Ed.) Artificial intelligence and the future of testing. Hillsdale, NJ: Erlbaum.

Ross, B. H. & Spalding, T. L. (1991). Some influences of instance comparisons in concept formation. In D. Fisher, M. Pazzani,& P. Langley (Eds.) Concept formation: Knowledge and experience in unsupervised learning. San Mateo, CA: Morgan Kaufmann.

Bradshaw, G. L., & Ross, B. H. (1993). Reminding and interpretation during encoding. Proceedings of the Fifteenth Annual Conference of the Cognitive Science Society, 249-253, Boulder, CO.

Blessing, S. B., & Ross, B. H. (1994). Content effect on problem categorization. Proceedings of the Sixteenth Annual Conference of the Cognitive Science Society, Atlanta, Ga.

Murphy, G. L., & Ross, B. H. (1994). Predictions from uncertain categorizations. Cognitive Psychology, 27, 148-193.

Ross, B. H., & Spalding, T. L. (1994) Concepts and categories. In R. Sternberg, (Ed.), Handbook of Perception and Cognition, Vol. 12. Thinking and Problem Solving. San Diego, CA: Academic Press, Inc.

Ross, B. H., & Bradshaw, G. L. (1994). Encoding effects of remindings. Memory & Cognition, 22, 591-605.

Spalding, T. L., & Ross, B. H. (1994). Comparison-based learning: Effects of comparing instances during category learning. Journal of Experimental Psychology: Learning, Memory, and Cognition, 20, 1251-1263..

Malt, B. C., Ross, B. H., & Murphy, G. L. (1995). Predicting features for members of natural categories when categorization is uncertain. Journal of Experimental Psychology: Learning, Memory, and Cognition.

Blessing, S. B., & Ross, B. H. (in press). The effect of problem content on categorization and problem solving. Journal of Experimental Psychology: Learning, Memory, and Cognition.

Ross, B.H. & Murphy, G.L. (in press). Category-based predictions: The influence of uncertainty and feature associations. Journal of Experimental Psychology: Learning, Memory, and Cognition.

Papers under review or revision

Ross, B.H. (under review-a). Classification and the effects of interacting with instances.

Ross, B.H. (under review-b). Category use affects classification.

Ross, B.H. & Kilbane, M.C. (under revision). Analogical mapping views and the effects of principle explanations.

GRADUATE STUDENTS INVOLVED IN THIS RESEARCH

Stephen Blessing (honors student, now graduate student at Carnegie-Mellon)
Matthew Kilbane
Susan Perkins
Martin Preslar
Thomas Spalding
Patricia Tenpenny

PRESENTATIONS BASED UPON AFOSR FUNDED WORK

Reminding-based learning, AFOSR Workshop on Cognition, Arlington, VA, November, 1989.

Some influences of instance comparisons in concept formation. Invited speaker, Workshop on Computational Approaches to Concept Formation, Palo Alto, CA, January, 1990.

Reminding-based problem solving and learning. Cognitive Science, Georgia Institute of Technology. Atlanta, GA, April, 1990.

Reminding-based generalizations in categorization. Midwestern Psychological Association, Chicago, May, 1990. (Spalding & Ross, presented by Spalding).

Reminding-based problem solving and learning. McMaster University. Hamilton, Ontario. December 1990.

Focussing effects of reminding-based category learning. Midwestern Psychological Association, Chicago, May 1991. (Spalding & Ross, presented by Spalding).

Conservatism in learning. GUV meeting, Bloomington IN, September, 1991.

Concept formation and the learning of feature correlations. (T. Spalding & B. Ross), MPA, May 1992.

Learning from the use of earlier examples: Evidence about what is learned. (M. Preslar & B. Ross), MPA, 1992.

Categorizing and solving algebra word problems. GUV meeting, Ann Arbor, MI, October 1992.

Instance comparisons during category learning. (T. Spalding & B. Ross), Psychonomics Society, St. Louis, November, 1992. [poster]

Reminding-based learning. University of Michigan, Ann Arbor, MI, December, 1992.

Remindings during learning. University of Chicago. Feb 1993.

Concept formation and the learning of feature correlations. (T. Spalding & B. Ross), MPA, Chicago, May, 1993.

Comparison-based category learning. (T.Spalding & B.Ross). Poster presented at the American Psychological Society Convention, Chicago, IL. June, 1993.

Encoding and remindings. GUV meeting, Ann Arbor, MI, October 1993.

The effects of remindings on understanding. University of Michigan, Ann Arbor, MI, October 1993.

Remindings in learning. University of Maryland, College Park, MD, March 1994.

The use of classified instances affects classification. Psychonomics Society, St. Louis, November 1994.

Category learning is not simply learning to classify: The effect of using categories on classification. Georgia Institute of Technology Cognitive Science Colloquium, January 1995.

Reminding-based learning. New York University Department of Psychology Colloquium, March, 1995.

DESCRIPTION OF RESEARCH

In this final report, the overview and background are given, followed by a summary of the results of each of the projects.

OBJECTIVES

When learning a new cognitive skill, novices spend much of their time solving problems. In doing so, it is common for novices to think back to an earlier problem that the current problem reminds them of and use this earlier problem to help solve the current problem (e.g., Ross, 1984, 1987, 1989a). This use of the earlier problem not only affects performance on the current problem, but also provides the learner with additional knowledge that can be accessed and used on later problems. The aim of this research was to understand the nature of the learning that results from this use of earlier problems. Little was known about this crucial source of learning.

Such within-domain analogies occur frequently during learning. In the view presented in this work, a generalization is formed from making an analogy between problems. Rather than positing a separate generalization process that operates upon completed instances, the generalization may be a byproduct of the analogy. In using the earlier problem to help solve the current problem, comparisons must be made and some aspects generalized over. Remindings, by setting up the analogy, may determine what pairs of problems are compared and, hence, what generalizations are made. The learning comes about because, while the noticing might be based on a variety of similar aspects (including superficial ones) between the problems, the comparisons forces the generalization of many of the aspects. My earlier work (e.g., Ross & Kennedy, 1990) showed that the use of earlier problems allows novices to begin to form generalizations across problems. Thus, this means of learning is one way in which novices can begin to develop more expert-like knowledge structures. The research conducted had three goals.

The first goal was to understand the nature of the resulting generalization, the information included in its representation. If the generalization results from the reminding, it is likely to be selective and conservative. Selective means that what is learned will depend upon how it is used. Conservative means that the generalization will be somewhat tied to the problems from which it arises. Some theories assume some selectivity and generalization, but little empirical work existed for necessary theoretical development. The work under this grant examined the specificity of what is learned, as well as its generality.

The second goal of this research was to examine the implications of these ideas for the development of expertise. Research on expertise in mathematics and physical science domains suggests that experts have problem schemas that allow them to categorize problems, as well as associated procedures for solving problems of that type (e.g., Chi, Feltovich, & Glaser, 1981). Despite the importance of these

schemas, little research examined how they are learned. A common idea is that they may develop from the comparison of problems. However, this idea leads to two questions? One, how do people know which problems to compare? Two, why are people comparing problems (i.e., what is the nature of this comparison)? The reminding-based learning view suggests that people compare problems when one makes them think back to another and they do so in order to use the earlier one to solve the current problem. In addition, this view suggests that people may develop problem schemas that are influenced by the superficial aspects, because these aspects are known to affect remindings. Thus, part of this project was concerned with the development of such problem schemas and the possibility that some schemas may critically depend upon superficial contents, even in experts.

The third goal of this project was to begin to extend this work to more everyday problem solving situations. This extension will not only allowed the application and test of these ideas in an important new setting, it also forced the extension of this work to important situations that are hard to experiment with in more formal domains.

PROJECTS

In this section, I provide details of the completed projects, organized by the three goals of this research.

1. Selectivity and conservatism of learning

These projects examined the representation of the resulting generalization. The focus here was on asking how the problem solving affects what gets incorporated into the generalization and the effect on later performance. Two projects are summarized. The first focused on how our use affects the problem categories we have. The second examined the interactions among the different processes involved in reminding, analogical transfer, and generalization.

The effect of problem use on the generalization.

A central idea of the reminding-based view is that the generalization depends upon the earlier problem comparison. This selectivity effect (that leads to conservatism) has been investigated in the context of category learning (Ross, Perkins, & Tenpenny, 1990), because in formal domains different earlier problems will often lead to very similar generalizations. This project investigated whether the details of how the problem is used (rather than which problem) affects what is learned. It can be viewed as a type of transfer-appropriate processing in problem solving. The project examined problem categorization as a function of use, both in terms of interacting with the instances and using the category to make further inferences.

Problem categorization as a function of problem use. This project integrates work on problem solving and work on categorization. The research in problem solving suggests that problem categorization is among the most crucial aspects of problem solving and one of the most difficult to learn. (One way to think about this is in terms of the applicability conditions used in production systems, which most computational modelers have argued is the toughest part of getting the systems to work). The research in categorization focuses on the instances being categorized as if the goal was to know the category rather than use the category to do something with the instance (e.g., make a prediction, inference, etc.). In this project, I took the problem solving perspective and showed that how the instances are used affects what is learned about the category. My first work in this area used simple algebra equations and examined how particular problem solving interactions with the instances affects classification. The second set of studies examined a case in which categories are used in a diagnosis task, again relating the use of the categories to the classification.

i. The effects of interactions with instances on classification (Ross, under review-a)

   (a) Contrasting categorization and use. I used two simple types of algebra equations that differ in terms of the order of the operations used to solve them. For example, a + (bx/c) = p can be solved for x by $\underline{s}$ubtracting a, $\underline{m}$ultiplying by c, and $\underline{d}$ividing by b, which I shall call SMD (for this experiment it does not matter if they go SDM instead). A second equation type, for example, (q+ mx)/b = s, can be solved by multiply, subtracting, then dividing (MSD). In addition, I confounded some mathematically irrelevant properties with each type, such as the letters likely to be used as constants (early in the alphabet vs. late) and the parentheses (whether it excluded part of the sum or not). Two groups learned to categorize these equations as Type 1 or Type 2. One group simply categorized them and were given feedback. The second group also had to solve the equations for x. The results show that both groups are excellent at categorizing new instances like the study ones, but the question of interest is what happens when the test items separate out the effects of operations, letters, and parentheses. The simple story is that the solving group is much more likely to choose on the basis of operations than is the categorization-alone group (.88 versus .70). Although this may not seem like a surprising result, the point to remember is that (1) these were very simple equations for which all subjects had much experience and the two types had obvious differences and (2) it demonstrates the importance of the use in learning the problem categories. The same results held with somewhat different problem categories, in which the coefficients could be combined simply (e.g., 4ax and 2ax) or required factoring out the x (e.g., 4ax and 2bx). The proportions of classifications based on the mathematical structure were .83 versus .59.

   (b) Different uses lead to different generalizations. One difficulty with these experiments is that it is hard to rule out a simple better learning explanation for the performance of the solving group. So, in the next experiment, two groups solved the problems differently and ended up with corresponding differences in their

understanding of the problem categories. In particular, I added in a y variable, such as in the equation $9 + sx/t = (5qy + b)/f$. One group solved for x (making it SMD), while the other group solved for y (making it MSD). The other types of equations were MSD for the x group and SMD for the y group. Thus, though they saw the same equations and categorized them in the same way at study (i.e., the equation given was called Type 1 by both groups), they were predicted to be learning very different things about what characterizes each type. This prediction was confirmed for 85% of the test items (in which the variable z was used). Even when predictive superficial characteristics were used at study and test, 68% of the responses were consistent with the particular use.

(c) Different uses affect category formation. Another experiment showed that the use of the instances affects what categories might be formed, as indexed by a sorting task. No category labels were given (i.e., it was an unsupervised learning task). One group solved for x (with the coefficient materials) while the other group copied the equations down and read them aloud (to ensure that they read them over carefully). After, all subjects were given the equations and asked to sort them into piles in terms of how they think about the equations. For the group that solved for x, 60% sorted by the coefficients, while the read-aloud group had only 10%.

(d) Different uses affect typicality. The use affects not only the categories formed and the classification of items, but also the internal structure of the category. The group solving the equations was less affected by superficial changes in rating the typicality of new instances compared to the group that simply categorized (a difference of 2.9 versus 3.6 on a seven point scale).

These experiments provide important findings on problem categorization for two reasons. First, they provide a way of examining more directly the selectivity issue, how the use of the problems affect what is learned from them, an idea of importance for both theoretical and instructional reasons. Second, they provide an important bridge between the huge amount of work on categorization to the much smaller body of work on problem categorization.

ii. The use of categories affects classification (Ross, under review-b)

The second set of studies investigated how the use of categories (as opposed to instances) may affect the classification. Using a common type of material in the classification literature, diagnosing diseases from a set of symptoms, I added a category use task. Subjects were shown the symptoms, but had to learn to classify patients by the diseases (i.e., a diagnosis task). This task has been often used and the predictiveness of the features is a major determinant of classification. However, after classifying into diseases, they then had to decide which treatment to give the patient (each disease had its own two treatments). Although several symptoms are presented equally frequently and are of equal predictiveness of the disease, two of the symptoms are also predictive of the treatment. The results show that these

treatment symptoms not only come to be viewed as more predictive of the treatment (which they are), but also as more predictive of the disease (which they are not): 96 versus .80. Following these basic results, the experiments presented in the paper show:

(a) The effects do not depend upon particular characteristics of the design. The same effect of category use is obtained whether the symptoms are perfectly predictive or probabilistically predictive (Experiment 2: .76 versus .69), whether the category use is intimately related to the classification (as treatment and disease are) or is totally arbitrary (as in having the use be the last names of patients, Experiment 3: .91 versus .79), or whether the classification task is introduced first and practiced alone before introducing the use (Experiment 4: .84 versus .73).

(b) People are learning not only about the symptom-to-disease connection, but rather the information from the use is more widely available to be useful for a variety of tasks. Experiment 5 uses a generation task, in which subjects were given the diseases after learning and had to generate symptoms that would occur with this disease. Treatment symptoms were more likely to be generated (.75 versus .53) and are, on average, generated earlier (average of second versus fourth). Experiment 6 simply asked how often each symptom occurred in the experiment, with no reference to diseases at all. Treatment symptoms were judged to have occurred more frequently (though they did not), 18.6 versus 16.2.

(c) The category use effect is influenced by prior knowledge people bring to bear. In this experiment, the treatment symptoms were related pre-experimentally (e.g., sneezing, itchy eyes, skin rash, which are all allergy-like symptoms) or the nontreatment symptoms were. Overall, the usual category use effect was obtained, but there was a strong interaction. When the treatment symptoms were related, they were classified much more readily than the nontreatment symptoms (.96 to .66). However, when the nontreatment symptoms were related, there was no difference between the treatment and nontreatment symptoms. Thus, prior knowledge was affecting the category use effect. When the two effects were "pointing" to the same symptoms, the effect was large. When the two effects were pointing to different symptoms, they canceled each other out.

Most theories in psychology and AI have focused on classification as an end in itself, when rather it is usually a means to accomplish some other goal. That is, in problem solving, one categorizes problems so that procedures for solving those problems can be applied. In diagnosis, the disease categories are used to help provide treatments or explanations. This project suggests that the uses made of the category feed back and affect the classification as well. Learning is tied to the use made of the category.

Interactions among different processes.

The work on selectivity can be viewed in a slightly different way - the

generalization depends upon the details of the problem solving. This idea led to a detailed consideration of how the different processes involved in reminding, application, and generalization may be interacting. This project was in collaboration with Prof. Gary Bradshaw, a faculty member at Illinois We focused on remindings and transfer. All views of analogical problem solving assumed the new (target) problem is represented and that this representation leads to a retrieval of some earlier problem (which is then mapped and transferred to the new target problem). However, there is also evidence that people may be reminded during the reading of a problem.

Our studies (Bradshaw & Ross, 1993; Ross & Bradshaw, 1994) examined whether such remindings of earlier problems may mean that the target problem's representation is affected by the reminding. The importance of this point is that much of the focus on mapping as a separate process may be misplaced, IF the earlier problem is affecting how the target problem is represented (i.e., much of the transfer occurs DURING the target being represented, not between the represented target and the earlier problem). So, if this problem makes me think of an earlier one, I may fit the current problem to the earlier one. Bradshaw and I tested this idea in a general way, using how simple stories are interpreted. Our results show that a simple cue (a proper name) reminds people of an earlier story and affects their interpretation of the current one (78% of the interpretations were consistent with the superficial cue). In addition, we found that this is an encoding effect in that one can see the effect on sentence-by-sentence reading times, with a slow-down of 316 msec for sentences inconsistent with the interpretation. The details of the methods and results are available in publication so are not given here. This work is important in showing that the processes interact and, in particular, that the reminding and representation are not separate, contrary to most current theories.

2. The development of problem solving expertise

As discussed earlier, reminding-based learning provides one perspective on the nature of problem schemas and their development. Some of the work already discussed can be viewed as very early precursors to problem schemas, but no evidence has yet been presented that in fact they do lead to problem schemas.

My focus has been on the specialized schemas that often include superficial information as well. This work is available publicly in a short form (Blessing & Ross, 1994) and will soon be coming out in a journal (Blessing and Ross, in press), so I give just a short description here. We wrote a number of algebra word problems (e.g., distance, interest, mixture), with each problem having three versions: one with appropriate (typical) contents, one with neutral contents, and one with the inappropriate contents, using the typical contents of a different problem type. For example, a motion problem (in which the typical instantiation includes vehicles moving at a certain speed for a certain time) would have appropriate contents of a vehicle such as a car, neutral contents of a nonvehicle such as an arrow, and

inappropriate contents of typical contents from a different schema such as interest accruing in a bank.

Our results showed that experts do use the superficial content *when it is predictive of the structure* (i.e., problem type). When highly experienced algebra problem solvers solve these problems, the inappropriate content leads to worse performance than the other two problem types (.63 versus .75 for the other two conditions). A verbal protocol study of subjects solving these problems shows that schemas are more likely to be accessed quickly with the appropriate contents.

In the second study, other highly experienced subjects received the same problems one clause at a time and were required after each clause to say what type of problem it was. Earlier work. showed people could do this quite readily, but we focused on differences between the appropriate, neutral, and inappropriate problems. We found that the appropriate problems were categorized faster than the neutral problems and the inappropriate problems were categorized slower than the neutral. The proportions of the problem needed to categorize were .29, .55, and .79 for the appropriate, neutral, and inappropriate problems, respectively. The inappropriate contents often lead to subjects initially making the (content-appropriate) inappropriate categorization, before "seeing" the right category. The appropriate problems were quickly classified, often on the basis of just a few words. The neutral problems took longer, but were almost always correctly classified by the end. One explanation for the lack of difference between the appropriate and neutral problems in the first experiment was that by the end of reading the problem, these relatively easy problems were classified so that schema information was available for solution in both the appropriate and neutral conditions.

The third experiment used more complex problems (with irrelevant information added) to further examine the appropriate and neutral conditions. When the problems were difficult to classify and solve, the appropriate problems were solved more readily (.66 versus .58). Together, these studies provide important information about the problem schemas and their use, particularly concerning the inclusion and use of superficial aspects. As a whole, these experiments provide the clearest evidence for the effect of problem content on the performance of experienced problem solvers. This idea is important not only for the study of expertise, but also as a clue as to how learning must proceed (i.e., the development of expertise). Predictive superficial contents are not abstracted over, but rather used as a very helpful heuristic as to the problem type.

3. Everyday problem solving

In addition to the work on probability theory and formal domains, I have examined remindings and inferences in less formal situations. Two projects were conducted under this idea. In one, I have investigated how remindings may affect the simple categories that are learned, resulting in reminding-based generalization. I view this work as parallel to the problem solving work, but it sometimes allows

me to more cleanly investigate certain issues, especially selectivity effects (Ross, Perkins, & Tenpenny, 1990). In addition, I have taken this work further (and relating it to the category use ideas mentioned earlier) to examine how categories are used in making predictions.

Reminding-based generalizations in category learning. The work on reminding-based generalizations in categorization was done with my graduate student, Tom Spalding, who is now an Assistant Professor at University of Iowa and has been published (Ross & Spalding, 1991; Spalding & Ross, 1994). Because this work is now available in publication form, I will just mention the main point. Within category learning, two views of representation are prototype and exemplar. The prototype view is that people abstract from instances and use that knowledge to classify new instances. The exemplar view is that people classify new instances by similarity to earlier instances. Much of the empirical results favor the exemplar view, but it also seems that we do have prototypes of categories and other abstractions. However, there are few ideas of how experience with exemplars leads to these abstractions. (The traditional view, that it happens automatically, is inconsistent with some results.)

Spalding and I argued that as people classify new instances by similarity to old instances, they may notice some commonalities between the instances and, from that, begin to abstract. Repeated uses of exemplars and earlier abstractions will lead to various abstractions that together can be used to classify new instances. Thus, remindings lead to abstractions as instances are compared. The papers provide a variety of evidence for these ideas. A final set of studies that are almost complete investigates how reminding may affect the *interpretation* of features. For example, if a person who rides a bike to work reminds one of someone who went swimming, the category may be viewed as more athletic, whereas if the earlier person installed a solar water heater, the category may be viewed as more energy conscious. Note that "rides a bike to work" is not an ambiguous action at one level, but the intention is ambiguous. It appears that remindings, by their selectivity may allow people to learn very different things from very similar information. This work has important implications for learning in complex informal domains, where many different things can be learned from each experience.

Predictions from uncertain categorizations. The second project is also investigating the use of categories, with a focus on how categories may be used for prediction. Although this is a little more tangential to the reminding-based learning work, it blends in nicely with the effects of category use on learning and also provides a more detailed assessment of how the categories are being used in a complex situation. This work is in collaboration with an Illinois faculty member, Gregory Murphy. We have taken the ideas about the use of categories and examined it more directly in how the categories are used to make judgments. In many situations in which people need to make a prediction or judgment about an object, it is assumed that they access the category information about the object and use this, along with any specific knowledge they might have about the object, to

make the prediction or judgment. However, often the categories are not certain, but rather people have only a likely category. An animal seen briefly is probably a cat, but it could be a dog or raccoon. How do people use this uncertainty in making predictions?

This project has led to three publications (Malt, Ross, & Murphy, 1995; Murphy & Ross, 1994; Ross & Murphy, in press). The short summary is that in many situations people essentially ignore the probabilistic nature of the categorization and act as if it were certain. That is, the likelihood of the alternative categories has little influence on the judgment, under a variety of situations that were used to try to find such an influence. This result suggests that categories are crucial in inferences because they determine what subset of knowledge will be used. We think this lack of influence of alternative categories derives from a heuristic organization of prediction brought about by basic information processing limitations (i.e., considering alternatives is too difficult in most situations). The first paper (Murphy & Ross, 1994) showed this lack of effect across a large number of different circumstances using artificial categories. We thought that the use of alternative categories in prediction might occur with more naturalistic situations (Malt, Ross, & Murphy, 1995), but we still found little influence. Our latest work (Ross & Murphy, in press) addressed more extensively the situations under which one might be able to find such an influence. If the alternatives are brought to mind at the time of prediction, either explicitly or by the nature of the prediction judgment, then some influence is found. This project has been important for my problem solving work as well, both in forcing me to think more concretely about how various categories may be used and also in considering the limits of information processing more directly in the learning process.

CONCLUSIONS

This project examined the common means of learning from remindings. A number of experiment sets have helped to provide a further understanding of these remindings and relate the work to other areas. The studies on the representation of the generalization found evidence for the selectivity and conservatism of generalizations and showed how such generalizations may be tied to the uses. This work relates to that of two important areas of research. First, given the importance of problem schemata for expertise (e.g., Chi et al, 1981), much research has examined schema acquisition. Remindings provide help to relate the problem solving directly to the learning and claims a specific mechanism is responsible. Second, the work on categorization has focused on the classification of instances as the goal. The experiments reported here extend the examination of categorization to the use of categories, within a problem solving perspective, and show how this can influence a variety of category-related processes, such as classification.

Another set of experiments examined how these remindings may provide a means of becoming more expert in a problem solving domain. Even highly experienced solvers rely upon superficial similarities that are predictive of the

problem type. The content had effects on problem classification, problem solving accuracy, and problem solving strategy. These results suggest, consistent with the reminding view, that predictive (but formally irrelevant) content may still be included for accessing the problem solving schemata underlying expert performance.

Third, an examination of remindings in everyday learning situations extended the findings and better tested some theoretical ideas. These studies explored the idea of how comparisons based on remindings may affect the representation. In addition, some experiments examined more directly how categories are used for inference in with both simple, artificial categories and with complex, natural categories.

The overall results of this project provide a clearer understanding of reminding-based learning and relate it to work on expertise, categorization, and schema acquisition.

References

(citations to work funded by this grant are included in the beginning of the report)

Chi, M. T. H., Feltovich, P. J., & Glaser, R. (1981). Categorization and representation of physics problems by experts and novices. Cognitive Science, 5, 121-152.

Hinsley, D. A., Hayes, J. R., & Simon, H. A. (1977). From words to equations: Meaning and representation in algebra word problems. In M. A. Just & P. A. Carpenter (Eds.), Cognitive processes in comprehension. Hillsdale, NJ: Erlbaum.

Ross, B. H. (1984). The effects of remindings on learning a cognitive skill. Cognitive Psychology, 16, 371-416.

Ross, B. H. (1987). This is like that: The use of earlier problems and the separation of similarity effects. Journal of Experimental Psychology: Learning, Memory, and Cognition, 13, 629-639.

Ross, B. H. (1989a). Remindings in learning and instruction. In. S. Vosniadou & A. Ortony (Eds.) Similarity and analogical reasoning. Cambridge: Cambridge University Press.

Ross, B. H. (1989b). Distinguishing types of superficial similarities: Different effects on the access and use of earlier problems. Journal of Experimental Psychology: Learning, Memory, and Cognition, 15, 456-468.

Ross, B. H. & Kennedy, P. T. (1990). Generalizing from the use of earlier examples in problem solving. Journal of Experimental Psychology: Learning, Memory, and Cognition, 16, 42-55.